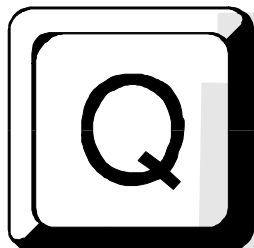

TPRS

*The official journal of the
leading regional amateur
radio digital communications
organization of the Americas*



Quarterly Report

November 1997

Published by the Texas Packet Radio Society · Volume 12, Issue 2

IN THIS ISSUE

President's Report 1

TPRS Fall Digital
Symposium 4

An Amateur 900 MHz
Spread-Spectrum Radio
Design (Part 2,
conclusion).....5

Telemetry in TexNet
(Part 2 of 3) 8

President's Report

Tom McDermott, N5EG

It's been a quick 3 months since the last Quarterly Report, and it's been a busy summer. First of all, TPRS will be holding it's annual Fall Digital Symposium on Saturday, December 6th in Austin (Fall lasts late into the year in Texas). Look for details on the location, program, etc. elsewhere in this issue. One of the things that we are trying to focus on in this symposium is to illustrate how things have changed in networking. It used to be that running TCP/IP meant setting up a NOS box (or later, a LINUX box), developing or finding custom drivers for TNC's, and setting up a lot of parameters. While these things are still possible, the new generation of operating system software (like Windows 95™ among others) includes TCP/IP built right into the OS, and they provide 10-base-T networking right out of the box.

Recently, I setup a 10-base-T LAN in my house, and it took a total of about 40 minutes to have it up and running between several computers. File and printer sharing is built in. A 10-base-T hub can be purchased

(Continued on page 2)

TPRS



Quarterly Report

The TPRS Quarterly Report is published by the Texas Packet Radio Society, Inc. Materials submitted for publication will be published on a space-available basis. Unless otherwise indicated, explicit permission is granted to reproduce any material appearing herein, provided credit is given to both the author and to TPRS.

Address articles and correspondence to:

Texas Packet Radio Society
P.O. Box 50238
Denton, Texas 76206-0238

TexNet Network Coordinator

Harry Ridenour, N0CCW @ K3WGF.#STX.TX.USA.NOAM

Texas Networks Group

Harry Ridenour, N0CCW @ K3WGF.#STX.TX.USA.NOAM

MailBox/BBS Group (Interim)

Dave Wolf, WO5H @ WO5H.#DFW.TX.USA.NOAM

Editor Q-Report

Brad Smith, KC5SP, brads@galstar.com

Database Manager

Frank Aguilar, N5SSH, faguilar@icsi.net

Texas Packet Radio Society

Web Site: <http://www.tprs.org>

Director (exp June 1999) /President

Tom McDermott, N5EG

Internet: n5eg@tapr.org

Director (exp June 1998)/Vice President

Joe Borovetz, WA5VMS

Internet: jsboro@intellex.com

Director (exp June 1999)/Secretary

Bob Morgan, WB5AOH

Internet: morganb@inetport.com

Director (exp June 1999)/Treasurer

Jim Neely, WA5LHS

Internet: jimneely@ibm.net

Director (exp June 1998)

Dave Wolf, WO5H

Internet: dwolf@tapr.org

(Continued from page 1)

for about \$60, and the appropriate 8-conductor cables are available off the shelf at many computer places. The 10-base-T interface board for my PC cost \$29 (less than many so-called 'hi-speed' serial interface cards!). You can even purchase "cross-over" 10-base-T cables which eliminates the need for a hub if you are just connecting two computers back to back. What does this have to do with packet radio? Several members of TPRS, together with TAPR have been developing a Spread Spectrum radio with a 10-base-T interface for TCP/IP radio. This radio is designed to support IP protocol in a plug-and-play fashion with these new Operating Systems.

Further, server packages are available for computers that are intended to be servers (like Windows NT™). Servers for EMAIL, Web sites, and other functions are available off the shelf. The spread spectrum radio permits using the radio to provide many of these functions for hams in a local area who want to get together and exchange email, web pages, photographs, audio files, etc. in real time at ISDN-like speeds. Other software applications are available, such as audio and video conferencing, but they are not yet quite so off the shelf (but close). What makes all this available to us hams is that the Spread-Spectrum radio looks exactly like an IP Internet connection to your computer, and thus you can use these off-the-shelf software packages without any modifications.

Setting up the networking involves setting IP addresses, and previously this has been difficult in amateur radio since you had to register with an address coordinator. However, the new radio will support a protocol called DHCP - Dynamic Host Control Protocol. What DHCP does is to dynamically assign an IP

(Continued on page 3)

(Continued from page 2)

address to your computer each time that you connect to the Spread Spectrum radio (actually the server has to assign you an address through the radio). This means that you don't need to have any address coordination in order to communicate using the latest applications. But you and your friends do need to coordinate your servers, EMAIL addresses, and that's why the ability to put POP3 servers on the system is important.

These topics will be discussed at the Fall symposium, and I think you will find it very informative.

1997 ARRL DCC Symposium a big success

The digital communications symposium was held in Baltimore, Md. In October, and the attendance was quite good. APRS was especially well represented, with an 8-hour seminar on all of the different applications possible with it. Also, GPS equipment was on display, along with homemade antennas and amplifiers. Tom Clark, W3IWI displayed his Totally Accurate Clock (TAC) based on the GPS receivers. There was one paper session on building a precision frequency reference by phase-locking a crystal oscillator to the GPS 1 pulse-per-second output. This would yield Rubidium-standard accuracy at a small fraction of the price.

There were also papers on Spread Spectrum, and the TAPR SS radio project in particular. Also, Dewayne Hendricks, WA8DZP, Phil Karn, KA9Q, and myself presented a 5-hour seminar on Spread Spectrum Sunday morning that was extremely well attended. Much of the material from my portion of the seminar will be presented at the TPRS Fall Symposium. It describes in some detail how the TAPR SS radio will provide all these IP services, and how the radio is designed.

You can also find the audio and slides from the seminar and the paper from the DCC on the TAPR web site, <http://www.tapr.org>. Many of the papers from the conference are available here, also. There was a really good paper on SNMP (that's the standard Internet Simple Network Management Protocol) by one of the student presenters on managing a TNC. We hope to extend this work to our SS radio, with the author's help, or the help of his advisor, a professor at the University of Sherbrooke.

Don't forget to bookmark the TPRS web site, <http://www.tprs.org>. It contains current information on TPRS activities, maps to the Fall symposium, etc. Dave Wolf, WO5H has done a marvelous upgrade to the page, you should take a look. Lastly, I would like to thank some of the volunteers who make TPRS, it's activities, and TexNet possible. We seldom think about how much work these folks contribute. Bob Morgan, WB5AOH, and Harry Ridenour, N0CCW operate and maintain the TexNet network daily. Frank Aguilar, N5SSH maintains our membership database, and Brad Smith, KC5SP puts together our newsletter. Jim Neely, WA5LHS, is our treasurer, and Greg Jones, WD5IVD is a valued advisor (plus he has helped us get our newsletter posted and mailed). -30-
(Windows 95 and NT are trademarks of Microsoft).

**TPRS Fall Digital Symposium
December 6, 1997
370 Sanchez Education Building
University of Texas, Austin**

Friday Evening - Informal get together:
5:00 PM Sholtz Garten

Saturday (December 6th) Morning:
8:30 AM - 12:00 AM SS radio hardware

1. Antennas, feedlines, propagation for 915 Mhz - Bob Morgan
2. Details on TAPR FHSS radio (repeat of DCC symposium) - Tom McDermott
3. FHSS radio design status update - Bob Stricklin
4. Update on Spread Spectrum Rulemaking - Greg Jones

Lunch Break 12:00 AM - 1:30 PM

Afternoon: 1:30 - 3:30 PM
TexNet Update

5. TexNet status update - Bob Morgan
Code revision
Analog monitoring, other neat stuff
6. TexNet Network status - Harry Ridenour
7. BBS / APRS status and issues - Dave Wolf

Late Afternoon: 3:30 - 5:00 PM
SS radio software

8. TCP/IP Services on modern platforms - Bill Reed
How IP works, routing, how TCP works
POP3, NEWS, Web servers TCP/IP on Windows 95, etc. How to install, connect to LANS
DHCP, permanent IP addresses, etc.
9. Open discussion - all

Conclude 5:00 PM

Texas Packet Radio Society, Inc.

TPRS was founded in 1985 and is an educational, public service, and scientific research non-profit corporation. Texas Packet Radio Society goals are:

- 1- design and research amateur radio packet networks
- 2- provide education in the area of general packet usage

To accomplish better communications in the region, TPRS has been organizing statewide working groups to cover various networking topics. The current working groups are the Mailbox/BBS Group, TCP/IP Group, and the TexNet Support Group. TPRS hopes that these working groups will help promote information exchange in their respected areas in Texas. New working groups are formed as needed to provide channels for discussion and to help provide direction for that area of digital communications. Anyone can participate in a working group; TPRS membership is not required.

TexNet

TPRS has established a digital packet network protocol, a standard hardware package for the network nodes, and software modules that implement the TexNet network.

The basic design philosophy of TexNet is an open, inexpensive, multi-resource, high speed 'backbone' with access through multi-connect capable local nodes. On the high speed side, TexNet is a 9600 baud network system. For local access, compatibility with the typical 2 meter AX.25, 1200 baud, AFSK/FM station is the operational norm. Other baud rates and modulation techniques can be supported on the primary user port or secondary port. The system is totally compatible with both versions of the AX.25 protocol specifications for user connections. With these general specifications, TexNet has been designed and tested to enable all users to take advantage of this high speed, full protocol protected packet network system.

Each node offers, in addition to TexNet access, local area digipeater service, 2 conference bridges for full protocol protected roundtable or net operation, a full multi-connect, multi-user mailbox system, a local console for installation and maintenance setups, a debugger module for long distance and local software monitoring, and an interface for a weather information server for regional weather information, if available.

The NCP-PC (TexNet for PC) creates a direct interface to the PC platform. The Z80 based PC card supports 4 channels for communications. This co-processor approach allows the AX.25 and TexNet-IP to run on the card without affecting the PC. This allows the full power of the PC to be used for network applications. The versatility of this board is only now being developed and applications are endless.

The TexNet Network

The Texas TexNet network system has been operational since October 1986. When fully operational, the network reaches from the border of Mexico to Missouri. Use of the Texas TexNet system is open to all amateur operators. TPRS has been coordinating the installation of the Texas TexNet system. Further expansion of the system depends entirely upon the amateur community.

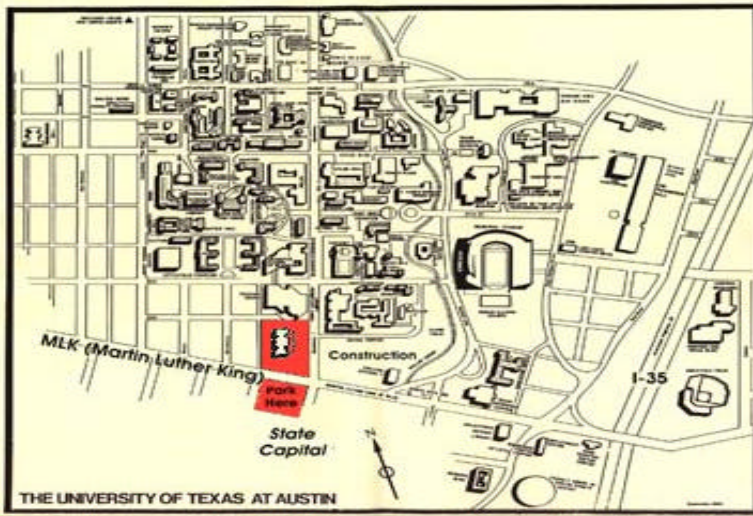
INFORMATION

TPRS is interested in spreading our information and research efforts as widely as possible. We want other groups involved with packet efforts to get in contact with us. We will provide information for those amateur packet groups that are interested in this system for their areas. If you would like more information concerning TPRS or TexNet, please drop a letter to:

**Texas Packet Radio Society, Inc.
P. O. Box 50238
Denton, Texas 76206-0238**

TPRS MEMBERSHIP

TPRS membership is widespread with most members located in Texas, but members are located in other states and in foreign countries. Membership is open to any interested person. If you are interested in becoming a member and receiving the TPRS Quarterly, please send your name, address and call with membership dues of \$12 per year. A membership application is available elsewhere in this issue.



**Map to the TPRS
Fall Digital
Symposium
at UT Austin**

An Amateur 900 MHz Spread-Spectrum Radio Design Part 2 of 2

Tom McDermott, N5EG,
n5eg@tapr.org,
Bob Stricklin, N5BRG,
n5brg@tapr.org,

Bill Reed, WD0ETZ, wd0etz@tapr.org

The use of FEC and QPSK provides at least 9 dB improvement in system gain as compared to uncoded non-orthogonal Frequency-Shift Keying (FSK) which is utilized in almost all commercial part-15 radios. However, the use of coherent modulation techniques increases both the cost of the radio and the difficulty of the design. We felt the 9 dB. performance improvement made this tradeoff worthwhile. Fortunately, Harris provides a DSP-based digital Costas-loop QPSK demodulator IC (the HSP 50210) which appears to have sufficient programmability to meet the synchronization speeds provided that some clever algorithms ("quick-lock") are employed.

Two risks are felt to represent the greatest challenges in the radio design. First is the ability of the hopping VCOS to settle to adequate frequency accuracy and stability within 10 milliseconds. Second is the ability of the Digital QPSK loop demodulator to achieve synchronization lock with our special "quick-lock" technique. The prototype design will be used to assess these design risks.

Block Diagram

Figure 1 is a block diagram of the baseband processing, processor, and LAN Interface portions of the radio. Figure 2 is a block diagram of the RF and IF processing parts of the radio. The radio design is based on a Motorola 68360 microprocessor. It controls all major functions of the radio, and the LAN interface. A Motorola 68160 provides the 10-base-T Ethernet port. FLASH memory is utilized solely in the processor, to allow updates of the code at a later time without physically opening the radio or removing / programming any EPROMS.

The data from the LAN port is buffered by the

(Continued on page 6)

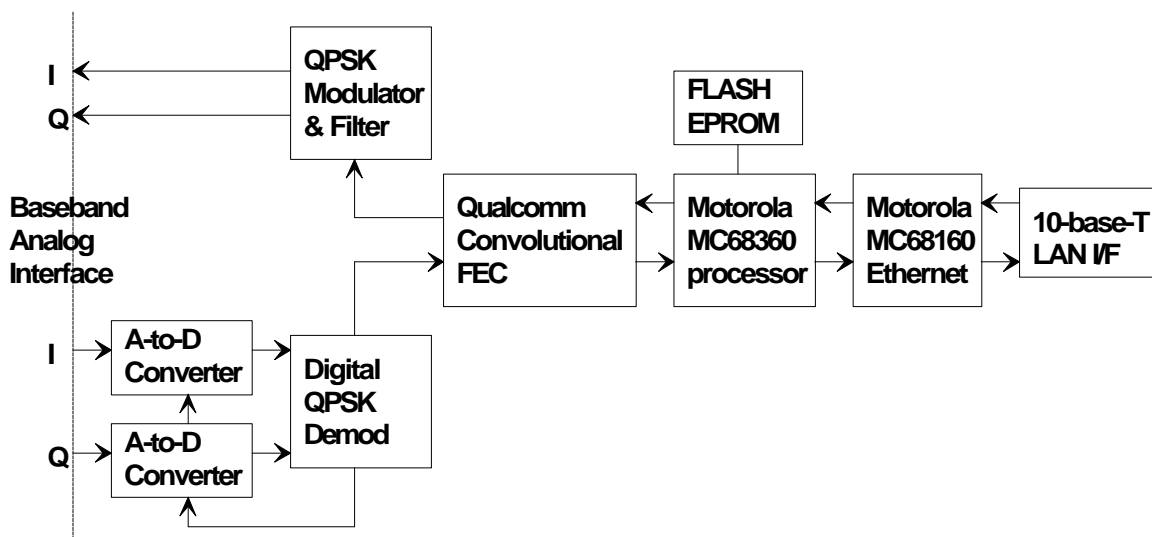


Figure 1 – Block Diagram: Baseband Processing and LAN Interface Circuit Description - Transmit Direction

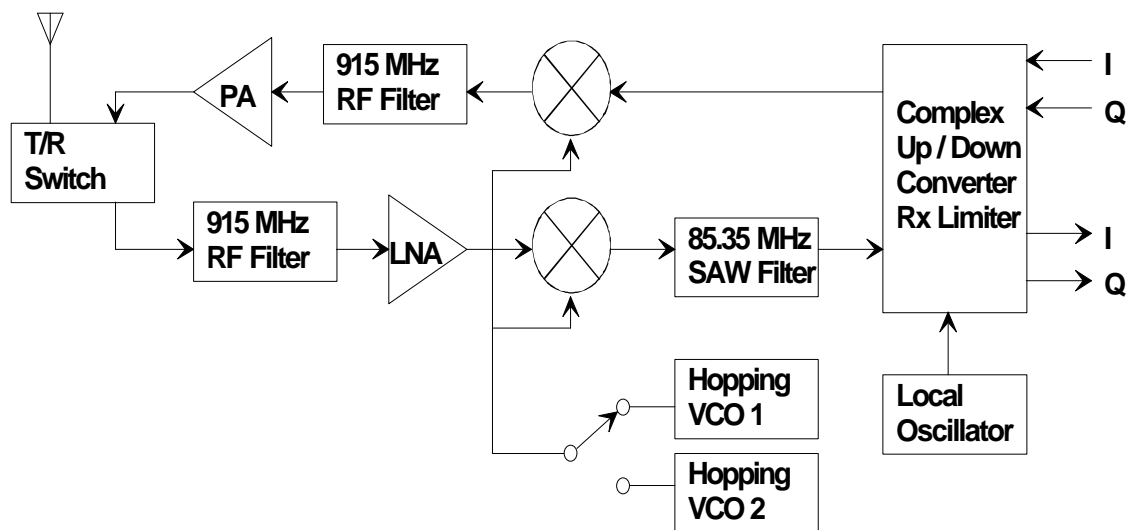
(Continued from page 5)

68360 and converted to a proprietary frame format based on HDLC and then sent to a Qualcomm convolutional coder IC. In modes 0 and 1, the coder produces two output bits for each input bit (rate = 1/2 mode). In mode 2, the code is punctured to rate = 7/8. These two bits become the in-phase (I-) and quadrature (Q-) channels to a Motorola QPSK modulator IC. The modulator IC provides raised-cosine roll-off at baseband of the two channels via an FIR filter. It also contains two D-to-A converters, and thus provides the I- and Q- analog baseband output signals.

The two baseband analog signal are connected to a Harris quadrature up-converter IC that generates I- and Q- signals at the IF frequency of 85.35 MHz. These signals are then further upconverted to the 902 MHz band, and filtered by a dielectric filter to elimi-

nate the IF image frequency. It is then amplified by a Motorola integrated PA chip to about 100 milliwatts. The signal is routed through a PIN diode switch and through a pair of directional couplers to the antenna connector. The directional coupler signals are rectified and filtered, and fed to an A-to-D converter chip. These signals provide measurement of the forward and reflected power levels.

In the receive direction, the signals are passed through a dielectric filter (to eliminate the image frequency) and then to a Motorola low-noise downconverter IC. From there they pass through an 85.5 MHz, 600 kHz wide SAW filter and an amplifier. At that point, they are sent to a Harris downconverter IC which provides a large amount of gain through a two-stage limiter, and then downconverts the signal to baseband, producing the I- and Q- baseband analog signals. These signals are



then digitized by a pair of 10-bit A-to-D converters, and sent to the Harris digital Costas-loop demodulator IC.

The demodulator IC first performs a complex frequency rotation to adjust for any frequency offset and phase error between the transmitter and receiver, then provides symbol timing and carrier frequency acquisition and tracking. Finally it provides AGC on the demodulated baseband signals, and performs a soft-decision threshold comparison of the I- and Q-channels against the reference level. These are in the form of two 3-bit words, one for the I-channel, and one for the Q-channel.

Motorola 68360. The microprocessor recovers and removes the HDLC frame, and transmits the received data out the 10-base-T LAN port via the 68160.

Hopping VCOS

The design utilizes two VCOS in a pair of phase-locked loops (PLLs). While one loop is operational on frequency, the other loop is busy slewing to a new frequency. At the end of each 10-millisecond period, the new VCO becomes the active VCO and the previously active VCO is slewed to another channel. In this manner, each VCO plays leapfrog, being utilized half the time. This allows each phase locked loop 10 milliseconds to achieve satisfactory frequency accuracy before it is switched into service.

All of the RF-determining reference frequen-

(Continued on page 8)

(Continued from page 7)

cies are derived from a single crystal-controlled oscillator. This oscillator is ovenized to minimize its error from the desired frequency during temperature excursions.

The actual programming of the VCO PLLS occurs by a small PIC chip (one-time programmable single chip processor). This chip contains the hopping sequence of the radios, and cannot be altered by the user. United States Department of Commerce regulations prohibit the export of FHSS radios from the United States if the hopping sequence can be altered by the user.

Synchronization

The most difficult part of any design is the synchronization of the transmitter and receiver, both in terms of the Transmit / Receive switching (T/R) and also in terms of carrier frequency acquisition. An initial synchronization interval occurs prior to the radios becoming linked. This takes some time to occur. The demodulator utilizes a sweeping process to recover carrier lock. However once this is achieved, the microprocessor is capable of reading out the frequency error at the receiver demodulator from the acquisition register in the demodulator. Based on the actual RF channel utilized during the initial synchronization, it computes the master-oscillator frequency difference between the transmitter and receiver. Subsequently, each time that the radio hops channels, the microprocessor computes the new effective frequency difference, and pre-loads the demodulator carrier recovery loop register with the proper frequency offset value to place the recovered carrier very close to the proper frequency. This helps the demodulator lock very quickly. This is the "quick-lock" technique referred to earlier.

Acknowledgments

We would like to thank the Tucson Amateur Packet Radio Corporation (TAPR), which is sponsoring this project.

See TABLE 1 on facing page.

Telemetry in TexNet

by Bob Morgan, WB5AOH, Austin
Part 2 of 3

There is also an 8 input version of this chip, the ADC0809, a 28-pin dip package, or the same size as the eeproms we use. Both are an 8-bit converter, which means that they can describe an input signal in terms of the nearest 8-bit number which corresponds most nearly to its actual voltage. In other words, it is an indication of the precision of the reading. An 8-bit number corresponds to a decimal number from 0 to 255 (unsigned) or a signed (a number with both plus and minus values, using one of the bits as a - sign) from -128 to +127. This particular application is using unsigned numbers, and they represent voltages on a linear scale from zero volts (common ground in the node) to the nominal +5V DC bus that drives the digital circuitry in the NCP or TNC2. I say nominal 5V because the common 5 volt regulators have some tolerance associated with their accuracy, and for digital circuitry they are commonly used with we might see voltages as low as 4.8 to maybe as high as 5.2 volts. I have seen some that were as high as 5.6 volts, although that is actually out of spec. Anyhow, we are stuck with whatever 5V bus we happen to have, and we have to record that value for later use, on all the readings for that particular node, when we do the arithmetic.

What is a "point" ?

A point is the common name for an input or output

(Continued on page 10)

Mode	End Points	Performance	Throughput
P P S	Point-to-point search	Search mode to establish initial link in PP mode.	
P P 0	Point-to-point (i.e. user end system to user end system)	Rate=1/2, half-duplex, 10 msec T then 10 msec R.	150 kb/s
P P 1	Point-to-point (i.e. user end system to user end system)	Rate=1/2, transmit slots as needed, communication of slot requests across link.	300 kb/s
P P 2	Point-to-point (i.e. user end system to user end system)	Rate=7/8, transmit slots as needed, communication of slot requests across link.	525 kb/s
P N S	Point-to-Control Link (i.e. user end system to control link of a multi-radio node)	Search mode to establish initial link to control channel of a node.	
P N 0	Point-to-Node (i.e. user end system to data channel of a node).	Rate=1/2, half-duplex, 10 msec T then 10 msec R.	150 kb/s
P N 1	Point-to-Node (i.e. user end system to data channel of a node).	Rate=1/2, transmit slots as needed, communication of slot requests across link, with node doing slot voting across all channels and downstream notification to all users.	300 kb/s
P N 2	Point-to-Node (i.e. user end system to data channel of a node).	Rate=7/8, transmit slots as needed, communication of slot requests across link, with node doing slot voting across all channels and downstream notification to all users.	525 kb/s

Table 1 – Proposed Operational Modes

(Continued from page 8)

channel of a data acquisition and control or telemetry system, and those systems may have as few as 4 or 8 points to several thousand for industrial control systems. They might be analog or discrete (digital on/off) types, of all conceivable ranges, or they might

be pulse counters that count events such as bottles on an assembly line or tons of coal on a conveyor belt. They could also be pulses from a rain gauge tipping bucket, and this use is within our reach. The points we have on these latest prototypes for TexNet have 16

analog inputs, scaled from zero to approximately 5 volts DC, and are single-ended inputs, where the other side of the circuit is assumed to be grounded. There are some things called double-ended or differential circuits that have two input terminals for both (+) and (-) sides of a line, and operate somewhat independently of ground, and measure the voltage difference between the two terminals. By its nature, the ADC0817 or 0809 is a single-ended device, having one input pin for each point desired, working against common ground for its (-) reference. They are extremely high input impedance CMOS devices, essentially they look like a small value capacitor to ground, with some protective junctions in parallel to Vcc and ground. I connect all input points through RC filters to remove noise and cut off higher frequency audio components.

How are the various "points" used on a node?

A TexNet NCP is a 3 radio channel system, although most nodes don't have radios on all three channels, or ports. We also use TNC2's in TexNet, and we can have either one or two radio channels on that device. There are typically four readings that are useful to have on a radio channel, and they are: limiter/S-meter reading, a frequency reading, an FM or FSK deviation reading, and some kind of a transmit power or drive measurement from the transmitter. I have reserved 4 inputs for each radio port, so that can use a total of 12 points of the 16 channel multiplexer on the converter, leaving 4 points available for other purposes. Since the +5V bus is not exact, and potentially could drift or change, I like to use one of the points for a fixed reference that can be used to

compare with at will to ensure accuracy, and typically I assign it to the last point in the multiplexer. Also there is usually at least one +12V power bus in the node, and sometimes there is a separate supply for the radios and for the NCP or TNC2, and it is useful to be able to read those remotely. It may also be useful or at least satisfy idle curiosity, to know the temperature of the rack or the room or the outside air, and a thermal probe can also be attached. We could actually attach a weather station or maybe a DF system. If you follow APRS, there are all kinds of potential applications here.

What is "scaling" ?

This particular converter operates between 0 and 5 V DC, or that is, it is "scaled" from 0 to 5 V. For instance, a thermometer may be scaled, or marked off, from -40 to +260 degrees or something similar. Our 8 bit converter scales the voltage from 0 to 5 V DC to a number from 0 to 255, and it is linear, or its "curve" is a straight line on a graph. We could mark off graph paper from 0 to 255 on the x-axis, and 0.00 to 5.00 volts on the Y-axis, and draw a straight diagonal line through zero, to the point where 5V and the number 255 intersect, and this would be our "curve", in this case straight, or linear. Or we can dispense with the graph paper, and just scale it with arithmetic. If we have a displayed reading of 062 decimal, then we divide 62 by 255, which is .243137 on a calculator (or 24.3137 % of scale), and then multiply that fraction by the 5 volt full scale reference, we then get 1.22 volts. I am using a 1.2 volt reference IC, a precision two terminal regulator, on Moody and some other nodes, and in fact this is the typical reading from it, and the one at Moody hasn't budged from this value in over a year.

We also have to scale points to fit the range of the converter. If we apply 5.2 volts DC to our 5 volt converter, it obviously is overranged, like a thermometer that goes to 120 degrees being immersed in 125 degree water. In our case it happens to just read its maximum value of 255, but doesn't tell us that is off scale. Some systems do tell us, but they cost more. We have to ensure that we can't go off scale. Also, we can't put 12 volts into our converter, that kind of overload would

damage it, like putting our 120 degree thermometer in boiling water or melting solder. The multiplexer in the converter also has its quirks, and one of them is that it is a CMOS switch. While it does have some protection against gross overload which may save it from damage from moderate overloads, a temporary overload can affect

adjacent channels, if it causes all of the gates to go into conduction at the same time, and I have seen that device behave that way, as well as others I have worked around. Since the scale has two ends, we can overload it in the negative direction also, and we have to prevent that too, if there is some possibility the measured circuitry could develop a negative voltage. Some can, and some inherently cannot.

So how do we measure a 12 volt bus, and it might soar to as high as 14 volts during typical regulation? We have to divide it down, or prescale it, to some range that will fit inside our 5 volt window. I typically use a 3:1 or 4:1 resistive divider for these tasks. Then we have to remember to multiply our reading we get after scaling to 5V, like we did above, by multiplying by 3, or 4, or whatever the scale value is. Some applications might require subtracting an offset, or maybe using a polynomial curvefit, or at worst case a look-up table for each value.

It depends on the nature of the measurement, and how its analog circuitry is laid out and scaled.

What is a point database?

We will eventually need a database to keep track of what is connected to each point at each node that we telemeter. Typically an analog database has an entry for each point in the system, and for a "distributed" system like our network, it would have 16 entries for each node. Each entry would have a name or description, a point address (a number from 0 to 15 which tells us which input signal we are talking about, and is its actual I/O address offset), which node in the network we are referring to, and some scaling information such as full scale voltage and any offset voltage. In addition we might have to describe a situation where the input is connected to something very non-linear in nature, like a test point voltage in a receiver IF or limiter stage, that can be related in a nonlinear fashion, to a signal strength at

the receiver antenna terminals. We could also include alarm information, if we wanted to be informed if the signal wandered out of its permissible range, into an area that we didn't want it to operate in. Unconnected, or left over inputs or outputs, are typically called spares. For instance, I can't buy a 12 input converter, just because I only need 12 inputs. I have to buy either an 8 or 16 input chip. It would be best to tie off the unused inputs to ground or an adjacent known voltage, otherwise they might eventually charge up and start to affect the multiplexers, which could degrade the readings that we do want.

What are sensors and transducers?

Sensors or transducers are devices that convert something we want to measure, to a voltage that is a function of what we are wanting to measure. I can make a very simple and linear temperature probe from 5 common diodes in series, that has one end grounded, and the other end fed from a constant current regulator which should be chosen as one that has extremely low variability with temperature. That is an example of a sensor. "Signal Conditioning" is a process of amplifying and rescaling a signal so that its range will fit in the range of our converter, and it may also include some noise filtering. For instance, our limiter measurement might only deliver 0.6 volts DC at saturation, and that may not give us much resolution. In fact, we would top out at a scale reading of 30, so we would have only 30 discrete measurement points on our scale, and that would throw away a lot of detail. We would do well to externally amplify that 0.6V signal to at least 4.2 or maybe 4.8 volts, with an op-amp, prior to sending it to the multiplexer. This scaling ratio or curve must be precisely known for that particular point, so that we know that we have to multiply by that number every time we read that point. Other examples of transducers might be: Pressure transmitters, RF power meters, barometers, flow meters, AC power transducers, strain gages, AC or RMS converters, pH cells, wind direction and speed sensors, humidity and dewpoint probes, chemical analyzers, vibration pickups, seismometers, tachometers, and in short, anything in the world that someone has a need to measure, that can be converted

(Continued on page 12)

(Continued from page 11)

to an electrical signal. They all have their scale ranges and some are nonlinear so they have some kind of a curve or curvefit that needs to be applied after their voltage output is read. Some transducers are read on a constant current basis, instead of a voltage basis, like 4.0 to 20.0 milliamps for instance. If you have a few hundred feet (or more) of wire in series with it, we don't want to read the voltage drop of the wiring, it represents error. If we have a constant-current output regulator in the transducer output, the current will be the same value regardless of how many feet of wire we have, as long as the regulator

has enough voltage headroom to actively regulate within. Lots of these transducers typically derive all of their internal operating power from this excess voltage, so we only need two wires and a 24 volt supply, and we place a precision 250 ohm resistor in the (-) lead of the transducer, and develop a 1 to 5V signal across this resistor, which we convert and read.

What is in the complete converter?

The prototype consists of a PC breadboard card about 3" x 4" or so, with a Z80 socket (and a pin header on the bottom side that plugs into the NCP or TNC2), the ADC0817 and socket, and 2 or 3 logic chips that drive it. We need a 74HC138 for an address and bus decoder, and 6 gates of a 74HC14, and if we don't want to bring in an external clock signal from someplace on the NCP or TNC2 clock chain, we can

add a 74HC393 counter (which we use half of), which will derive our needed clock from the clock that is available driving the Z80. We need also one more external lead for a completion interrupt, it tells the NCP or TNC2 that a conversion cycle has been finished, and the number needs to be read from the converter. We typically connect it to the CTS-A pin of the first SIO chip, since the CTS pins can service interrupts, and we aren't using them for anything. The 74HC138 decodes both read and write I/O address accesses in the range of 30h to 3Fh (or 30h to 38h if we use the ADC0809), and this is an I/O address space that we don't use for anything on either the NCP or the TNC2 in TexNet. We directly gate this into the address select line of the

ADC0817/0809 multiplexer, which we also supply with the first 4 (or 3) address lines of the bus, and the write pulse selects, or turns on, one of our input pins so that it may be read by the converter. It requires a few microseconds for the multiplexer to turn on the selected signal path.

There is a RC time delay filter between a couple of the gates of the 74HC14, and another gate fires a delayed version of the write pulse at the "start" pin of the converter. This time delay gives the signal from the multiplexer time to settle to an accurate version of our desired input. Conversion to a number begins with the start pulse. The clock we feed to the converter causes successive approximation of each bit of the 8 bit converter, one bit per clock cycle. We attach to a 614 KHz source in the TNC2, or to a 307 KHz source in the NCP, or we can divide the NCP Z80 clock to 500 KHz with the 74HC393. Any of these clocks will work, and are in the optimum range of the converters capability.

The faster the clock, the faster the conversion. The converter is a CMOS charge comparator, and it initially charges internal on-chip capacitance cells to our input voltage, and then disconnects them and rearranges them one at a time in binary values of capacitance value, and compares thresholds, and comes up with each bit, in order of less significance. When 8 cycles have been completed, we have converted all 8 bits, and they are sitting in the output register. Then the converter fires the completion interrupt at our NCP or TNC2, and the software takes over. It looks up the address of the input that it knows is being converted (it just started it a fraction of a second earlier), and reads that address from the bus, which causes another gate of the 74HC138 and 74HC14 to fire, and that gate commands the converter to place its 8 bit value of the databus, which is connected to the

Z80 databus on the same circuitboard. The software then stores that number away and clears the memory flag to indicate that the converter is free to be used for some other input conversion, and returns to whatever it was doing when the interrupt arrived. As far as the rest of the software in the NCP or TNC2, in other words TexNet, it is set up to read and convert 4 values of inputs (according to which radio

the packet is being sent or received on) during each packet. It also periodically scans each input on a

regular schedule, so that all the values of all the inputs can be looked at on demand, for instance the environment variables, 12V bus, the reference, etc. It also gives us a random snapshot of what the radio values are, and we can usually see the idling no-signal values present quite often. Finally there are remote and local network command and response functions so that the data can be inquired from afar.

The higher layer software in the NCP or TNC2 provides some inquiry functions of data. Some of the lower level drivers in the software handle the packet I/O on the serial radio channels, and these drivers have been modified to select and start the converters when the address field of a packet is being sent or received. It is assumed that most radio channels are set up half duplex, but that the actual radio either is, or can be made capable of, receiving its own transmitted signal. We routinely read all of the 3 receiver parameter addresses during both receive and transmit. Additionally during transmit, we read the 4th value assigned to the transmitter. Most modern solid state transmitters have some kind of variable power driving stage, controlled by a voltage regulator which is actively controlling transmit drive, or at least protecting against bad VSWR, overheating, overcurrent, or similar gremlins. Some transmitters might also have RF power bridges which can be read. Anyhow, some significant point in the transmitter can be brought out and represent the health, if not the power output, of the transmitter. It is primarily provided for long term trending, and troubleshooting.

Since we can read our own transmitter frequency and deviation through our receiver (assuming we leave the receiver on during transmit), we can compare our own frequency to all of the other remote stations that could be connected, as well as compare one received signal against any other one. This gives us some more flexibility in troubleshooting, since it might be easier to get maintenance access to one node, and then infer where all of the other nodes frequency and deviation are running by reading the inaccessible nodes values. The limiter will of course be in hard saturation during transmit (and typically may even fold back

some), and trending this value over time is useful. When setting up a station's telemetry, it is necessary to "calibrate" the limiter, frequency and deviation signals against a known RF analyzer. The limiter curves are highly nonlinear, and require lookup tables (I haven't yet found any suitable curvefits for those, and I have tried). For instance the lookup table for Moody's limiters of course there are now two radios, it has a user port now) give numbers from 063 to about 192 from the trunk radio limiter, and yield a range of signal strength from roughly 0.1 uV (or -20 dBu) to the saturation point at around 6.0 uV (or 15.6 dBu). Note that while it converts numbers for each received packet, it starts converting during the address field of the packet, and the packet may or may not be successfully received errorfree. If it is not an error-free packet, the radio data is thrown out with the bad packet itself. We can only report the radio parameters associated with error free packets, and at present we only track connected stations, we don't track other packets from stations we aren't connected to. I have seen signal values on the Moody trunk radio with signal readings of as low as .25 uV, and that is not very much signal at 9600b UHF FSK, but it wouldn't be displayed if that packet hadn't been error free. But signals that weak don't perform very well. It takes about .75 uV or so to make a robust link, one that doesn't retry on long packets. That .25 uV reading was probably a poll, or maybe the short request for the analog reading itself, it probably wasn't a long I-frame of BBS output or something like that. Considering that site is located quite a way up on a TV transmitting tower, in a high RF environment where I have heard horror stories of non-successful use of DVM's with test leads picking up RF, that is pretty good performance. In fact, the resting or idle limiter reading of Moody's trunk receiver is about 035 on the converter scale, or probably about 0.05 uV, so I don't think we are coupling much noise into our receiver at that site. If intermittently we do have large amounts of noise that take out packets, this method won't see them since it only gives us readings of packets successfully copied. Indications of performance of that particular site tells us that most of its difficulties when it has trouble handling data probably are due

(Continued on page 14)

(Continued from page 13)

to it hearing too many of our own nodes during band openings, and when it is hearing signals, AX.25 is not allowing it to transmit, because it of course waits for the frequency to become clear before transmitting. absorption.

What is trending? How does one make use of the data?

Trending is the accumulation and comparison of data over time. We will have some broad range of readings for a process, most of which can be considered normal, as opposed to abnormal or even process alarm conditions. If we have no clear-cut issues that say which ranges are OK and which ones aren't OK, trending can become useful. If a process is running acceptably, its readings or internal values are probably acceptable also, particularly if they don't change over long term time, or at least repeat with repeating conditions or comparable repeating environments. Unexplained changes sometimes allow us to spot trouble brewing, and alert us to look around at the whole process to see what may be a problem in the making. Sometimes fixed limits are used, and that is pretty typical. Less often it may be necessary to use calculations like ratios or time rates of change instead of the actual values to spot trends. Some times it is necessary to use statistical techniques, and I have seen cases of deterioration of heat exchangers that could only be seen over several years of time, not just year to year changes.

How does this apply to parameters in a typical packet node?

We probably would measure bus voltages and maybe internal temperature. Sometimes the temperature is controlled, and sometimes it isn't. It might alert us to replace a cabinet fan on the next service trip on one node, or call a site host at some other node to inform him of a suspicion that his air conditioner or heater has failed. If it isn't a controlled environment, it just tells us what the heat stresses might be, or if it is a wide open site like Moody's TV tower cabinet, it follows the ambient temperature on the tower. If the wind isn't blowing, a tower can get hot enough in the late afternoon as to be downright hard to grip with your hands, and I have seen this effect this summer. A trend of no

changes of anything from year to year might indicate that we just let Moody node run another year without any attention (from March through July 1997 this was indeed the decision made). If we have a battery floating on charge, we may have tight limits on charging voltage, and they may vary slightly with temperature, and we will want to check, maybe even log them over time. As for radio parameters, we know we have some average rates of frequency drift on UHF crystal controlled radios. The numbers I have seen may be as much as 500 Hz downward per year. We have temperature stabilized some crystal holders, and this seems to help some. Certainly any abrupt change in receiver sensitivity, or no signal noise level is cause for concern. We can look at the normal daily cycle of variation of signal strengths of paths, and spot off normal propagation. For instance in the last week or so of July this year, we experienced degraded UHF propagation from Waco to San Antonio, and it was reflected as slightly lower than expected signal levels at Moody. During this time frame, we had a dome of high pressure overhead, and as soon as a front brought a few storms through and cleared it out, our paths returned to normal in August.

What are the commands? Which nodes can issue them? Which ones are local, and which ones are remote?

CONTINUED NEXT ISSUE

QR Submissions to:

Brad Smith KC5SP at
brads@galstar.com or
TPRS address on page
2.



TPRS Membership Application

Name _____ Callsign _____

Address _____

Apt. or Mailstop _____

City/State _____

Zip _____ E-mail address _____

Evening Phone () _____ Work Phone () _____

Membership is \$12 per year. How many years are you paying for? _____

☐ New Member ☐ Renewal

Make check payable to: Texas Packet Radio Society

FOLD HERE

PLACE
STAMP
HERE

ATTN:
Membership

Texas Packet Radio Society
P.O. Box 50238
Denton, Texas 76206-0238

Texas Packet Radio Society, Inc.

“Packet Radio - The Next Generation” is coming!

**Find out what it is
at the TPRS Fall Digital
Symposium, December
6th, 1997 at the UT
Austin campus.**

**Maps inside, or visit the
TPRS web page.**

Be sure to visit the TPRS web page:

<http://www.tprs.org>

**for the latest information on TPRS
activities.**

**A current listing of Packet nodes,
frequencies, and networks is located in the
North American Digital Systems
Directory (NADSD) on-line at:
<http://www.tapr.org/directory/index.html>**

Texas Packet Radio Society
P.O. Box 50238
Denton, Texas 76206-0238

ADDRESS CORRECTION REQUESTED

**BULK RATE
U.S. POSTAGE
PAID
DENTON, TX
PERMIT NO. 325**

To: